

09/641,223 PTO-892

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization International Bureau



(43) International Publication Date
30 September 2004 (30.09.2004)

PCT

(10) International Publication Number
WO 2004/084092 A1

(51) International Patent Classification⁷: G06F 17/30 (72) Inventors: CHOI, Arthur; 9 Milestone Manor Court, Germantown, MD 20876 (US). LEYBA, Todd; 12013 Evening Ride Drive, Potomac, MD 20854 (US). PORST, Beate; 2815 La Terrace Circle, San Jose, CA 95123 (US). SOMANI, Amit, Radheshyam; 7051 Quai Cove Ct., San Jose, CA 95120 (US).

(21) International Application Number: PCT/EP2004/002380

(22) International Filing Date: 13 February 2004 (13.02.2004)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data: 10/394,842 21 March 2003 (21.03.2003) US

(71) Applicant (for all designated States except US): INTERNATIONAL BUSINESS MACHINES CORPORATION [US/US]; New Orchard Road, Armonk, NY 10504 (US).

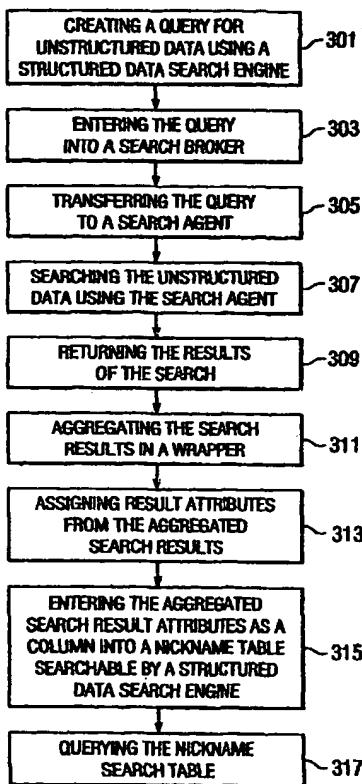
(71) Applicant (for MC only): COMPAGNIE IBM FRANCE [FR/FR]; Tour Descartes, 2, avenue Gambetta, F-92066 Paris La Defense Cedex (FR).

(74) Agent: ETORRE, Ives, Nicolas; Compagnie IBM France, Direction de la Propriété Intellectuelle, F-06610 La Gaude (FR).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM,

[Continued on next page]

(54) Title: REAL-TIME AGGREGATION OF UNSTRUCTURED DATA INTO STRUCTURED DATA FOR SQL PROCESSING BY A RELATIONAL DATABASE ENGINE



(57) Abstract: Method, system, and program product for searching unstructured data. The unstructured data, as text data or image data, is searched by first searching the unstructured data through an extended search broker. The extended search broker has intermediary components between a search requestor and a search agent for unstructured data searching. The search agent returns the search results to the broker for aggregation. The broker aggregates the search results and returns the search results to a wrapper. The wrapper then takes result attributes from the aggregated search results and makes the attributes available, e.g., as one or more columns in a nickname table. This nickname table is searchable by a relational database using a structured query language.

WO 2004/084092 A1



TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) **Designated States (unless otherwise indicated, for every kind of regional protection available):** ARIPO (BW, GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Published:

— with international search report

REAL-TIME AGGREGATION OF UNSTRUCTURED DATA
INTO STRUCTURED DATA FOR SQL PROCESSING BY A RELATIONAL
DATABASE ENGINE

5

FIELD OF THE INVENTION

Our invention relates to methods of searching (i.e., querying) unstructured data stored in a database, including converting an SQL query into a proper query for unstructured data.

10 A further aspect of our invention is translating an external SQL access of the type identified with structured relational databases into an internal access to unstructured databases or files, and translation of the SQL external query format into an intermediate or internal query format. The method, system, and program product of our invention is especially applicable to text searching and indexing.

15

BACKGROUND OF THE INVENTION

Accessing structured data with SQL is quite different from the full text search of unstructured data such as documents on the web. Structured data in the relational model

20 is maintained in two-dimensional tables as rows and columns. Each row in a table represents an instance of an object while each column represents the attributes of the object. A column is given a symbolic name and is assigned a specific data type (such as integer, date, etc.). Integrity constraints can be applied to columns to further indicate valid values.

25

Because column values are named and represented in a consistent format, you can select rows very precisely, based on their contents. This capability is especially helpful in dealing with numeric data. You can join together data from different tables based on matching column values. You can do useful types of analysis, such as listing objects in

30 one table that are missing from a related table (or that are present in a related table, or that

have specific attributes). You can extract specific rows of interest from a large table, regroup them, and generate simple statistics on them.

By contrast, unstructured data is not always organized in a consistent and predictable 5 manner. Unstructured data is stored in a variety of shapes and forms, distributed throughout the enterprise, and managed by the most appropriate software for the task at hand. The data tends to be recorded in free text form (for example, text contained in e-mails, notes, and documents) with little or no metadata codified into fields. As a consequence, searching is less parametric and more keyword-based in nature. Search 10 results derive more from what "matches" a given set of keywords than from computational criteria.

Yet it is desirable to query unstructured data in a structured way to add still more value to the results set. It would be advantageous to treat the web as a relational database, one that 15 could be queried using standard SQL. Just as importantly, it would also be advantageous to be able to treat a plurality of heterogeneous and unstructured data sources uniformly through an SQL interface, thus removing the ambiguity of their integration.

A conventional approach to solving these problems would be to extract the desired data 20 from the unstructured data sources, apply any necessary conversions to the data and then place the so-converted data into a relational database for later processing. Indeed this warehousing approach is a common method used today for a variety of applications.

However, this approach does not address the overarching issue of making unstructured 25 data available for parametric querying.

SUMMARY OF THE INVENTION

We address these problems by a method, system, and program product for parametric 30 searching unstructured data. The unstructured data, as text data or image data, is searched by first searching the unstructured data through an Extended Search broker. The extended

search broker has intermediary components between a search requestor and a search agent for unstructured data searching. The search agent returns the search results to the broker for aggregation. The broker aggregates the search results and returns them to the wrapper. The wrapper then takes attributes from the aggregated search result and makes 5 them available as one or more columns in a nickname table. This nickname table is searchable by a relational database using a structured query language (SQL).

One aspect of our invention is a method of searching unstructured data. The method comprises searching the unstructured data through an extended search broker. The 10 extended search broker has intermediary components between a search requestor and a search agent for unstructured data searching. The search agent returns the search results from the search agent to the broker for aggregating the search results to a wrapper. Next, the result attributes are entered into a nickname table, where the nickname table is searchable by a relational database using a structured query language.

15 A further aspect of our invention is a computer system having an unstructured data search system, and a structured data search engine, logically or virtually joined by a nickname table. The computer system is adapted to initiate a search of unstructured data initiated with the structured data search engine. The unstructured data search system includes an 20 extended search broker and a search agent for unstructured data searching. The unstructured data search system receives search results from the search agent and sends the search results and search results attributes to the extended search broker for aggregation into a wrapper. The computer system is adapted to enter the wrapper attributes into the nickname table. The structured data search engine is adapted to conduct 25 searches of search results attributes in the nickname table.

30 A still further aspect of our invention is a computer program product comprising computer readable code to program and configure a computer system to search unstructured data through an unstructured data search system that includes an extended search broker and a search agent for unstructured data searching. The program product further includes code capable of causing search results and search result attributes from

the search agent to be returned to the broker for aggregation, and to aggregate the search results and search result attributes in a wrapper. The code then enters the search result attributes in a nickname table, that is, as one or columns in a nickname table, where the nickname table is searchable by a relational database using a structured query language to
5 search unstructured data.

According to our invention unstructured data may be searched through the combination of a federated relational database engine and an extended search engine, exemplified by the IBM DB2 Federated Database and the IBM Lotus Extended Search, respectively.

10 Together the combination of a federated relational database engine and an extended search engine provide a relational interface to unstructured data that is distributed throughout the enterprise.

15 RDBMS federation provides an extensible architecture by which both internal and external developers can write wrappers to integrate external data sources. A federated RDBMS wrapper is the module that encapsulates (or *wraps*) the necessary native API calls to retrieve data from the foreign server. The wrapper is activated through callbacks at various points in the SQL processing cycle when data in its associated nickname table is referenced for processing the query over the remote data sources.

20 A nickname table is a special kind of RDBMS table used specifically for federation. It does not contain persistent data but rather is populated with data by the wrapper on demand. In this context the nickname table is a virtual table – not really existing in the physical sense but still supporting most if not all table operations that can be expressed in
25 an SQL expression.

30 An extended search engine, such as IBM Lotus Extended Search, is a product that “extends” the standard capabilities of other search products by allowing a user to search and retrieve all types of information, not just data that is managed by the hosting software. An extended search engine is based on a brokered search technology that provides for efficient, parallel searching of thousands of data sources distributed

throughout the enterprise. Extended search engine agents perform remote searches in the native grammars of the backend database systems. The agents return results to the broker for aggregation and delivery to the original caller – in this case the extended search engine wrapper.

5

By integrating these two technologies as described herein, several benefits are realized. First, unstructured data from a diverse set of data source types is made available in real time to the RDBMS engine through loading, including just-in-time loading, of the extended search engine nickname tables. By this expedient, data is now current and up-to-date. While other federated wrappers provide current data for a single source, only the extended search engine wrapper enables this capability for multiple data sources at a time. In this way the data received by the extended search engine wrapper is presented as a single normalized data set even though multiple sources were involved in its formation.

10

15 A second benefit is that the data remains distributed throughout the enterprise logically or virtually closest to where the work is performed. The extended search engine, and not the relational database engine, is responsible for the searching of thousands of data sources stored in different locations. The extended search wrapper used in conjunction with the extended search broker acts as a gateway to this vast amount of unstructured information.

20

20 A third benefit is the ability to apply relational concepts against the aggregate of unstructured information. With an extended search engine alone it is not possible to relate results received from one source with results received from other discrete sources in order to derive a subordinate results set. For example it is not possible to join together data 25 from different sources based on matching field values. Extended search engines only perform aggregation (analogous to a UNION query in SQL) on the returned data to produce a single result set. But, once the data is made available to the nickname table through the extended search engine wrapper, the data is free to be joined by the RDBMS engine with any other table in the database, including other nickname tables.

30

This capability has profound implications. By configuring a nickname table for each of the different types of data that can be retrieved by an extended search engine using concepts of our invention, it is now possible to use the power of the relational engine to easily integrate that data outside of its hosting environment.

5

The Figures

Certain embodiments and exemplifications of the invention are illustrated in the FIGURES appended hereto.

10

FIGURE 1 illustrates the various hardware and software platforms, which may be utilized in the practice of the invention. Illustrated are two end user clients, one connected to a data server through the World Wide Web and a web server, and the other connected to the data server through a local area network. The web server is connected to two relational database management systems, and to an extended search server, which is in turn connected to two unstructured databases.

15

FIGURE 2 illustrates the various elements of the method and system of the invention, with a user entering an SQL query to a relational database management system. The relational database management system is a federated relational database management system, and searches unstructured data using an extended search system with an extended search broker and agents. The results are accessed through an extended search nickname table and an extended search database.

20

FIGURE 3 illustrates a method of carrying out the query process by creating a query for unstructured data using a structured data search engine. The structured data search engine enters a query into an extended search broker. The extended search broker then transfers the query to a search agent, which searches the unstructured data, and returns the results of the search, aggregating the search results in a wrapper. Results attributes are assigned 30 to the wrapper, and the attributes are entered into a nickname table that is searchable by a

structured data search engine. The structured data search engine then queries the nickname search table.

DETAIL DESCRIPTION

5

Our invention facilitates searching unstructured data through a relational database management system using a structured query language. The invention comprises searching unstructured data using a structured query language and a relational data base management system through an extended search broker. The extended search broker has 10 intermediary components between a search requestor and a search agent for unstructured data searching. The search agent returns its search results to the extended search broker for aggregating the search results in a wrapper having structured query language searchable attributes. The wrapper attributes are entered into a nickname table, where the nickname table is searchable by a relational database using a structured query language.

15

A further aspect of our invention is a computer system having at least two search engines, one an unstructured data search engine, and the other a structured data search engine. The search engines are logically, that is, virtually, joined by a nickname table. The computer system is adapted to conduct searches of unstructured data through the synergistic 20 interaction through a nickname table of the structured data search engine and the unstructured data search engine.

25

A still further aspect of our invention is a computer program product comprising computer readable code to program and configure a computer system to search unstructured data from a relational database management system for structured data through an unstructured data search system including an extended search broker that has intermediary components between a search requestor and a search agent for unstructured data searching. The program product further includes code capable of causing search results from the search agent to be returned to the broker for aggregation, and to 30 aggregate the search results and search result attributes in a wrapper. The code then

enters the result attributes in a nickname table, where the nickname table is searchable by a relational database using a structured query language to search unstructured data.

5 The Integrated System of A Federated Relational Database Management System and an Extended Search System

According to our invention unstructured data is searched through the combination and interaction of a federated relational database engine and an extended search engine.

10 exemplified, respectively, by the IBM DB2 Federated Database and the IBM Lotus Extended Search. Together the combination of a federated relational database engine and an extended search engine provide a relational interface to unstructured data that may be distributed throughout the enterprise.

15 RDBMS federation provides an extensible architecture by which both internal and external developers can write wrappers to integrate external data sources. A federated RDBMS wrapper is the module that encapsulates (that is, wraps) the necessary native API calls to retrieve data from the foreign server. The wrapper is activated through callbacks at strategic points in the SQL processing cycle when data in its associated nickname table is referenced.

20

A nickname table is a special kind of RDBMS table used specifically for database federation. It does not contain persistent data but rather is populated with data by the wrapper on demand. In this context the nickname table is a virtual table – not really existing in the physical sense but still supporting most if not all table operations that can be expressed in an SQL expression.

An extended search engine, such as IBM Lotus Extended Search, is a product that “extends” the standard capabilities of other search products by allowing a user to search and retrieve all types of information, not just data that is managed by the hosting software. An extended search engine is based on a brokered search technology that

provides for efficient, parallel searching of thousands of data sources distributed throughout the enterprise. Extended search engine agents perform remote searches in the native grammars of the backend database systems. The agents return results to the broker for aggregation and delivery to the original caller – in this case the extended search
5 engine wrapper.

Certain embodiments and exemplifications of the invention are illustrated in the FIGURES appended hereto.

10 FIGURE 1 is a diagram illustrating the various hardware and software platforms, which may be utilized in one exemplification of the practice of the invention. Illustrated are two end user clients, 101 and 102, one, 101, connected to a data server, 121, through the World Wide Web, 105, and a web server, 111, and the other connected to the data server, 121, through a local area network. The web server, 111, is connected to two relational
15 database management systems, 131, 132, and to an extended search server, 141, which is in turn connected to two unstructured database servers, 151 and 152. The unstructured database server, 151 or 152, may contain one or more unstructured databases, and may alternatively, be web clients, where still further web server-data server combinations contain the unstructured data to be searched.

20 FIGURE 2 is a diagram illustrating the various elements of the method and system of the invention, with a user, 201, entering an SQL query, 202, to a relational database management system, 203. The relational database management system, 203, is a federated relational database management system, and using an unstructured data search
25 system with an extended search broker, 204, and agents, 205 and 205', searches unstructured data, 206 and 206'. The results are stored in an extended search nickname table, 207, and an extended search database, 208.

30 FIGURE 3 illustrates a method of carrying out the query process by creating a query for unstructured data using a structured data search engine, 301. The structured data search engine enters the query into the unstructured search system including the extended search

broker, 303. The extended search broker then transfers the query to a search agent, 305, which searches the unstructured data, 307, and returns the results and the result attributes of the search, 309, aggregating the search results and result attributes in a wrapper, 311. Result attributes are assigned to the wrapper, 313, and the results attributes are entered 5 into a nickname table that is searchable by a structured data search engine, 315. The results attributes are entered into a nickname table as a column therein. The structured data search engine then queries the nickname search table, 317, and returns the query results to the requestor.

10 By integrating federated database and extended search technologies as described herein, several benefits are realized. First, unstructured data from a diverse set of data source types is made available in real time to the RDBMS engine through loading, including just-in-time loading, of the extended search engine nickname tables. Data presented to the structured database search engine (the relational database management system) is now 15 current and up-to-date. While other federated wrappers provide current data for a single source, only the extended search engine wrapper enables this capability for multiple data sources at a time. In this way the data received by the extended search engine wrapper is presented as a single normalized data set even though multiple sources were involved in its formation.

20 A second benefit is that the data remains distributed throughout the enterprise logically closest to where the work is performed. The extended search engine, and not the relational database engine, is responsible for the searching of thousands of data sources stored in different locations. The extended search wrapper used in conjunction with the 25 extended search broker acts as a gateway to this vast amount of unstructured information.

30 A third benefit is the ability to apply relational concepts against the aggregate of unstructured information. With an extended search engine alone it is not possible to relate results received from one source with other discrete sources in order to derive a subordinate results set. For example it is not possible to join together data from different sources based on matching field values. Extended search engines only perform

aggregation on the returned data to produce a single result set. But, once the data is made available to the nickname table through the extended search engine wrapper, the data is free to be joined by the RDBMS engine with any other table in the database, including other nickname tables.

5

This capability has profound implications. By configuring a nickname table for each of the different types of data that can be retrieved by extended search engines using concepts of our invention, it is now possible to use the power of the relational engine to easily integrate that data outside of its hosting software.

10

Illustrative Example

By way of example, suppose an intelligence agency is trying to identify individuals who 15 might be potential terrorists. Given that this agency has access to a diverse set of sensitive sources, they may want to pose the following question: "List the names of those persons who have obtained a visa within a certain time frame, have purchased large quantities of fertilizer since those dates, have corresponded through e-mail with other individuals about bomb fabrication, and have applied for a class B (truck) drivers license."

20

The sources that need to be searched are many and have varying degrees of structure. Each source, however, has the notion of a person (identified by name) that can be used to perform the join. An extended search system would be used to perform the brokered searches to the remote data sources and return the appropriate data but would then rely on 25 the relational engine to perform the join. In order to do the join, the data needs to be organized into separate nickname tables – one for each type of data source being searched.

The example above requires four nickname tables. One nickname table would be 30 configured to represent e-mail results, another nickname table would represent drivers license information, another would represent fertilizer purchases, and another visa

applications. Each nickname table is associated with one or more extended search sources that are similar in purpose. For example, queries applied to the drivers license nickname table would, in reality, cause an extended search of fifty Motor Vehicle Administration data sources (one for each state).

5

It is assumed that each nickname table would contain a foreign key for the person's name on which the join could be performed. Extended search field mapping would be used to normalize the foreign keys in the nickname tables with the syntactical differences encountered in the backend sources (more on this later). Given that the four nickname 10 tables are configured in this fashion it would now be possible to pose the following query in SQL:

Table: VISA

FirstName	LastName	Date	CountryFrom
Bill	Stalwart	09/05/2001	Germany
Mary	Boutreaux	07/21/2001	France
Ali	Mohammed	09/10/2001	Iraq

15

Table: Driver

FirstName	LastName	Date	Clas	State
Bill	Stalwart	09/05/2001	A	MD
Mary	Boutreaux	07/21/2001	A	NY
Ali	Mohammed	09/10/2001	B	CA

20 **Table: Mail**

From	To	Date	Body
Bill Stalwart	Sally Cruthers	09/05/2001	Text of email

Don Crition	Martin Jones	09/05/2001	Text of email
Ali Mohammed	Mark Billsman	09/05/2001	Text of email

Table: Purchases

FirstName	LastName	Date	Description	Amount
Bill	Stalwart	09/05/2001	Refrigerator	\$673.00
Ali	Mohammed	09/05/2001	Fertilizer	\$2230.00

```

SELECT      N1.FirstName, N1.LastName
5   FROM      INS      as N1,
              DRIVER    as N2,
              EMAIL     as N3,
              FERTILIZER as N4
WHERE        ES_SEARCH(N1.DOC_RANK, 'DATE_OF_ISSUE >=
10  "01/01/2000") = 1 AND
                  ES_SEARCH(N2.DOC_RANK, 'LICENSETYPE = "B"')      = 1
AND
                  ES_SEARCH(N3.DOC_RANK, 'DOCUMENT TOKEN "bomb" and
                                      DOCUMENT TOKEN "fabrication"') = 1
15  AND
                  ES_SEARCH(N4.DOC_RANK, 'ProductType = "FERTILIZER" and
                                      Quantity > 500 and
                                      Date > "01/01/2000"')      = 1
AND
20  N1.LastName = N2.LastName AND
      N2.LastName = N3.LastName AND
      N3.LastName = N4.LastName

```

As previously mentioned Extended Search has the ability to map a common field name to one or more semantically equivalent but syntactically different fields on the backend. This can be very useful in producing a meaningful data model from the nickname tables.

5 For example, the symbolic identifier for person name in the fifty drivers license databases would most certainly be different – especially since the states are free to build their own MVA databases. Through Extended Search a single field for licensee name could be defined and mapped to the appropriate native fields in the drivers license databases. The single mapped field for licensee name would then be used in the drivers license nickname table. Without this mapping feature, you would be forced to define a nickname column for each uniquely defined name field in the backend database. The table would grow horizontally with many columns and would not be in 3rd normal form. The table would be sparsely populated since any one row would have only one name column filled in at a time.

10

15 There are applications however that can make use of a large number of fields coming from many discrete data sources. For these applications the relationships among the data (particularly unstructured data) are not known ahead of time thus making it very difficult to define and structure the nickname tables. Extended Search can support this type of application by populating what is referred to as a *vertical* nickname table.

20

A vertical nickname table contains four predefined columns. They are: field name, field value, field type, and source name. Rather than using mapped fields as recommended in the previous example, you do the reverse by letting Extended Search return the native field name/value pairs. Meta data such as the type of field (e.g. date, integer, etc...) and the source from where it came is also returned. Now the nickname table is densely populated with many rows of data – one row for each field per data source.

25 But how might these vertical nickname tables be used in the relational model? Suppose that the same intelligence agency were to ask a series of questions that helped them discover useful relationships among the data regardless of where it came from. They have

set up an extended search system to query over a thousand data sources of varying types each with their own data models and sets of fields. Not knowing which fields will be relevant in your search you configure extended search to search and return all of the fields. These fields are returned in the vertical format just described.

5 Now the intelligence agency can use the relational engine for the discovery and analysis of the data relationships themselves. They might ask for all fields that contain the word "fertilizer" ordering the result by data source name.

10

Federated Relational Database Management Systems

A federated database management system, also referred to as a federated database system or federated system is a system where each database server is an autonomous and independent centralized DBMS that has its own local users, local transactions, database

15 administrator, and, therefore, a high degree of local autonomy. In a federated database management system, each individual server can authorize access to specific portions of its database by specifying an "export schema" which specifies portions of its database that can be accessed by specific sets of non-local users, and the specifications that these non-local users have. In a federated database management system, a user is essentially an
20 additional interface to several local database management systems, thereby allowing a global user to access multiple local, autonomous databases.

A federated database management system, such as IBM DB2 Universal Database federated system, supports applications and users submitting SQL statements referencing

25 two or more DBMSs or databases in a single statement. An example is a join between tables in two different DB2 databases. This type of statement is called a distributed request.

A DB2 Universal Database federated system provides support for distributed requests across databases and DBMSs. A user can, for example, perform a UNION operation
30 between a DB2 table and an Oracle view.

A federated database management system, such as an IBM DB2 federated system, provides location transparency for database objects. If information (in tables and views) is moved, references to that information (called nicknames) can be updated without any 5 changes to applications that request the information. A further aspect of a federated database management system is compensation for data sources that do not support all of the particular SQL dialects, artifacts, or certain optimization capabilities of each other. Operations that cannot be performed under such a DBMS (for example, a file system that may not be able to implement a GROUP BY clause) are run under DB2.

10

A federated system functions in a semi-autonomous manner, that is, IBM DB2 queries containing references to Oracle objects can be submitted while Oracle applications are accessing the same server. A federated system does not monopolize or restrict access (beyond integrity and locking constraints) to other objects on the remote data sources.

15

In the case of the IBM DB2 Universal Data Base federated system, the system consists of the IBM DB2 UDB instance, a database that will serve as the federated database, and one or more data sources. The federated database contains catalog entries identifying data sources and their characteristics. A data source consists of a DBMS and its data.

20

Applications connect to the federated database just like any other database.

25

Federated database catalog entries contain information about data source objects: what they are called, what information they contain, and conditions under which they can be used. Because the federated database catalog stores information about objects in many data sources, it is called a global catalog. Object attributes are stored in the catalog. The actual data sources being referenced, modules used to communicate with the data source, and DBMS data objects (such as tables) that will be accessed are outside of the database. In this regard it should be noted that the one exception is a federated database can be a data source for the federated system. Specifically, IBM DB2 supports federation over 30 remote data sources and remote DBMS's.

Federated objects can be created using, for example, the IBM DB2 UDB Control Center, or SQL Data Definition Language ("DDL") statements. As a general rule, the required federated database objects are wrappers, servers, and nicknames. Wrappers identify the modules (DLL, library, and so on) used to access a particular class or category of data source. Servers, in this context, define data sources. Server data includes the wrapper name, server name, server type, server version, authorization information, and server options. Nicknames are identifiers stored in the federated database that reference specific data source objects (tables, aliases, views). Applications reference nicknames in queries just like they reference tables and views.

10

After a federated system is set up, the information in data sources can be accessed as though it were in one large database. Users and applications send queries to one federated database, which then retrieves data from DB2 family and Oracle systems as needed. User and applications specify nicknames in queries; these nicknames provide references to tables and views located in data sources. From an end-user perspective, nicknames are similar to aliases.

Many federated system operate under some restrictions. For example, distributed requests may be limited to read-only operations. In some federated systems utility operations (LOAD, REORG, REORGCHK, IMPORT, RUNSTATS, and so on) cannot be executed against nicknames. However, a requestor can use a pass-through facility to submit DDL (Data Definition Language) and DML (Data Manipulation Language) statements directly to database managers using the SQL dialect associated with that data source.

25 Federated systems tolerate parallel environments. Performance gains are limited by the extent to which a federated database query can be semantically broken down into local object (table, view) references and nickname references. Requests for nickname data are processed sequentially; local objects can be processed in parallel. For example, given the query

30 `SELECT * FROM A, B, C, D,`

where A and B are local tables, and C and D are nicknames referencing tables at Oracle data sources, one possible plan would join tables A and B with a parallel join. The results are then joined sequentially with nicknames C and D.

5

Extended Search Engines and Extended Search Applications

An extended search engine "extends" the standard capabilities of other search products by allowing an end user to search and retrieve all types of information, not just data that is 10 managed by the hosting software. For example, using a standard web browser, an end user can easily and quickly locate and review information contained in thousands of data repositories that are spread throughout the organization. These repositories may be of varied content and structure, and they may be geographically dispersed throughout the world.

15

With a single query, the end user can search internal and external web sites, full text indexes, Microsoft indices, document management systems, e-mail systems, file systems, relational databases, LDAP directories, and the full complement of Lotus databases - simultaneously. An extended search engine can be thought of as a brokered agent 20 working on the user's behalf: it searches each of the data stores in their native form, and returns results in a single aggregated results set.

The extended search engine shields an end user from data source diversity. This is because an end user interacts with a single, generally function-rich, extended search 25 engine interface that transparently distributes the end user's search to potentially thousands of diverse and disparate data stores. Results can be presented in a variety of ways.

With an extended search, an end user experiences more than what the end user might 30 have come to expect from a typical search application. An end user discovers that his or

her reach has been extended further into the enterprise, and that he or she can access information not easily attainable before.

Almost all search systems involve the use of an index that catalogs the information to be 5 searched. Without an index, the time to perform the search would be greatly exaggerated - much like trying to find a book in a library without a card catalog. But while most search solutions require an end user to re-index the information into a new index, an extended search engine, such as IBM Lotus Extended Search, leverages an enterprise's current data management investment.

10

Note that enterprise information usually exists in many shapes and forms, is distributed throughout the enterprise, and is generally managed by the task specific software for the task at hand. An extended search engine, such as IBM Lotus Extended Search, taps into this maze of data sources and modes of data access, and treats the existing collection as 15 one virtual index with a single point of entry.

Regardless of the mode of entry, an extended search engine translates the query into the native search languages of the target data sources and uses the search and retrieval methods that are native to each data source to return results.

20

One thing to recognize is that backend data sources are not created equal, however, and they can have totally divergent architectures, methods, schema, metadata and capabilities (such as full text indexes vs. relational databases vs. e-mail systems). Extended search tools and engines attempt to minimize these disparities whenever possible and achieve 25 more than the least common denominator effect.

For example, in querying multiple data sources, an extended search engine, such as IBM Lotus Extended Search, might combine two or more operations against one data source to achieve what can be done in a single operation on a different source. This approach 30 enables users to search multiple data sources, of differing types, architectures, structures, capabilities, and formats, in parallel with a single query. By this expedient it is possible

to maintain the mechanisms of the native database or search engine and preserve the integrity and currency of the information, while avoiding the duplication of dispersed data into a centralized index (reducing overall storage requirements and eliminating costs and overhead of re-indexing resources. Moreover, the data remains in close (logical and 5 virtual) proximity to the work that is being performed, and the respective data specific and task specific data sources and search engines can perform search operations in parallel, thereby improving search response times.

10 The extended search engine paradigm also offers greater scalability than multiple data source paradigms built around a centralized index. As the number of documents in a domain grows, so does the time it takes to index the information. It is not uncommon for the indexing of large domains to make the search impractical.

15 In this situation, the power of an extended search engine complements depth of the index: first, stabilizes the index at its maximum capacity, and then searches the index in conjunction with other data stores. With this approach, an end user have the option of indexing only what needs to be indexed while using the extended search engine to search both indexed and non-indexed sources.

20 Connecting new sources to the search domain results in little overhead being incurred by the extended search engine, and it is insignificant when compared to the expense of the actual search operations themselves. Because the search is executed on one or more machines that already exist in a network (including an internet or intranet) using existing search operations, the incremental cost of connecting new sources is negligible.

25 The extended search engine implements a strategy for accessing all of an enterprise's information wherever it happens to be located by using the enterprise's existing data management software. With this strategy comes the inherent fact that not all backend data sources are created equal. There are general areas in which most systems vary with 30 respect to search and retrieval. Building on these commonalities, an extended search engine offers a common ground that enables system implementers to work

complementarily within these areas of (1) Search language, that is, how the query is expressed; (2) Data model support, that is, how the data is organized, related, and represented, and (3) Application program interface (API) support, that is, how information access is performed.

5

With respect to the search language, nearly all data management systems employ a grammar or query language of some kind to express the criteria of a search. These grammars can vary widely depending on the structure and composition of the data. For example, in a free text system such as the web, for example, the search is generally expressed as a list of keywords. Additional notations are used to express boolean conditions (and, or, not) or positional information, such as specific words that must occur within the same sentence or paragraph. By way of comparison and contrast, if the data is highly codified and structured, the grammar may be more parametric and may support fielded operations (for example, the value of the Quantity field is greater than 100).

10

Clearly, it is impractical for a user to know the syntax of the entire set of search grammars used in an extended search. It is much more practical to let the user express the query in a single common search language. In IBM Lotus Extended Search the common language is the Generalized Query Language (GQL). It may be treated as a superset of search grammars from which most queries can be expressed.

20

All IBM Lotus Extended Search queries are expressed internally in GQL format (with the exception of "pass through" queries which will be discussed shortly). When the GQL expression arrives at a specific data source, Extended Search transforms it into the native query language of that source. This translation process maps the lexical elements of GQL to equivalent elements of the native grammar.

25

While the fullness of the GQL grammar is supported by all components of IBM Lotus Extended Search, not all backend search engines are capable of supporting all elements of the GQL grammar. When an extended search engine, such as IBM Lotus Extended Search, encounters these disparities, it tries to compensate by combining two or more

native operations to achieve the same effect. Only after the extended search engine has exhausted all alternative methods will it ignore the non-supported portion of the expression. Under these circumstances, an end user can instead pass through a query that does not get translated.

5

Common Data Model

Just as search grammars can vary with each dissimilar backend system, so can the data models used to organize and store information. The data model used by a particular data management system is typically designed for the class of applications it serves. This

10 determines the amount of structure and granularity found in its information.

For example, free text systems tend to use a loosely structured document model with low data granularity. A document may consist of a few fields (such as title, author, and body) but its text remains free in form and unstructured. By comparison, information can be

15 highly structured, such as that found in relational databases. Here, data is organized into rows and columns that can be related in any number of ways, which results in high data granularity.

A modern extended search engine, such as IBM Lotus Extended Search, normalizes this

20 diverse set of data models into a single model, one that is easy to understand and usable by most search applications. The search application designer need only contend with this one conceptual model of the data and not be confounded by the many.

Typically a source maps to a database in the conventional sense but can just as easily be

25 mapped to a web site, a directory in a file system, or a node in an LDAP hierarchy.

Likewise, a document maps conveniently to those systems that are text based but could just as easily represent an instance of data, such as a row in the table of a relational database.

30 A common problem encountered when relating data sources of different types is the mismatch in field labels. For example, an author's name might be labeled AUTH_NAME

in one data source and CREATOR in another, and yet be represented as three fields (such first name, middle initial, and last name) in another.

An important feature of the common data model is the ability to define mapped fields. A
5 mapped field is a composite of one or more native fields. To resolve the ambiguity in the previous example, an end user could define a single mapped field with the label AUTHOR. An end user could then map this field to one or more native fields in each of the data sources that support the semantic of author's name.

10 The benefits of this feature are compelling when used in the search expression. A user needs only to specify the mapped field in the query. IBM Lotus Extended Search will automatically associate the mapped field to the correct native fields on the backend. This approach greatly simplifies the query expression, and provides greater benefit as the number of data sources increases.

15 To illustrate this feature, consider the case of an application that is to display a barometer that is calibrated to reflect the age of a document (that is, more recent documents would have a higher barometer reading). If the "date created" field were labeled differently in each of the backend sources then the application would need an exhaustive set of
20 conditional statements, one for each discrete date field name. But, with mapped fields, only one mapped "date created" field needs to be identified for retrieval and subsequently tested. By referencing this single mapped field, the result would contain the correct date value as it correlates to the source's native date value.

25 On the other hand, if the result came from an e-mail system, an end user might want to return the date, subject, and author. These semantically different data fields can be selectively retrieved using their native source field identifiers. When the native field label is too cryptic for display purposes (for example, \$Doc_Abstract), then a mapped field could again be used as an alias for the native field name (for example, Document
30 Abstract).

A further aspect of an extended search engine is the provision for common API's. The common API's answer the question of "How do I interact with all of these different backend systems to search and retrieve their managed information?" The API to each backend system can vary dramatically by calling methodology, syntax, semantics, and

5 programming language. The functions issued through an extended search engine common interface are translated automatically by the system into the native methods of the backend system - much like the process of translating GQL into the native search grammar.

10 Extended search engine seek to interface with each of the different mechanisms on the backend in the best way possible to achieve optimum parity between all data sources while respecting the operating policies of each native search API. By using the published API of the hosting software, the integrity and security of the data source is preserved.

15 To achieve this communicability, a link module is created for each type of backend data source and is used to encapsulate all of the native API calls required for that type.

An agent-based technology is used to apply these link modules to their respective data sources. Agents work on behalf of a brokered search and are generally transparent (links

20 and agents will be discussed in more detail later on). From a designer's perspective, an end user is interfacing with one searchable backend system through a single API.

In the case of the IBM Lotus Extended Search, the common API is available in two forms. The first and most flexible is a Java bean interface. With this method the beans are

25 combined with the power of the Java programming language to develop a broad range of simple to highly complex and specialized search and retrieval applications. The product demonstrates the use of these beans in a sample web application that uses JavaServer Pages (JSPs).

30 Alternatively, a set of HTML-like tags that allow an end user to embed extended search engine functionality into new or existing web pages may be used. These easy to use tags

empower a web master and enable users to enter queries that specify a variety of search options. An end user can embed extended search engine tags anywhere within a web page - they do not interfere with the surrounding HTML tags.

5 The search results from an extended search application can be presented in a variety of ways. As part of configuring the system, the system can be programmed and configured to determine which fields a user is allowed to query, view on the results page, and retrieve from the data source. The user is presented with a single, consolidated page containing results from multiple sources. The list is pre-pruned for relevancy, thus
10 ensuring that the user sees the best matches first.

Extended search systems are characterized by scalability. Specifically, the distributed component architecture of an extended search system, such as IBM Lotus Extended Search, offers the flexibility to scale a system according to changing requirements. It also
15 allows the extended search engine components to be arranged in a topology that matches the environment, enabling an end user to blend IBM AIX, Sun Solaris, Windows 2000, and Windows NT platforms as needed.

Exemplary scalable extended search architectures include:

20 Vertically, within a single extended search server, an end user can configure multiple instances of server processes to influence the number of simultaneous requests that the server can process.

25 Horizontally, with multiple machines, an end user can set up additional extended search servers and additional web servers. For each extended search server, an end user can determine the types of server tasks an end user want to run. By having multiple servers, an end user can distribute and balance the processing load.

30 Extended Search Architecture

The huge economic growth experienced toward the end of the Twentieth Century has been attributed to the advances made in information technology and to those companies that invested in it. Because the technology was (and still is) changing at break neck speeds, companies frequently invested and then re-invested their Information Technology

5 ("IT") budgets into ever evolving products to manage their information.

More often than not, the result has been islands of information distributed throughout the organization - highly specialized for the task at hand, but not easily accessible on an enterprise wide basis. Even if a company were able to dictate a common IT policy for the enterprise, it could be thwarted by the single acquisition of a company that employs a

10 different IT architecture.

An extended search engine of the type herein contemplated, and exemplified by IBM Lotus Extended Search, has a multi-tiered design. For example, the IBM Lotus Extended Search system employs a four-tiered architecture. Messages start from search applications

15 in the first tier and proceed consecutively through subsequent tiers to the backend. In most cases, the backend is a third-party data source to which IBM Lotus Extended Search is connected but it can also be the Extended Search configuration database (CDB), a private backend that is managed by a relational database management system, such as IBM DB2 RDBMS.

20

Message flows between the tiers can be divided into two basic categories:

Run Time messages, which are messages usually issued by the user community to perform searches and retrieve documents.

25

Administrative messages which are issued by the Administrator and result in updates to the configuration database.

30

Links are the software modules that encapsulate the native API calls for search and retrieval to a specific data management system. They contain all of the required data

structures, programming objects, and procedural logic necessary to interface with the backend data system.

A link module is uniquely assembled to support (at a minimum) callable methods that

5 typically exist in all data management systems:

Methods to connect to and disconnect from the host system

Methods to search content and retrieve data from the system

10

The link module performs a null operation for those methods that are not supported by the backend source. For example, a file system search does not support the concept of connecting and disconnecting.

15

Translators are the software modules responsible for translating the incoming GQL expression into the native search grammar of the backend data system. They, too, contain all of the required data structures, programming objects, and parsing logic necessary to generate a syntactically correct search expression.

20

In some cases, the same translator module applies to several different backend systems, as is the case for the SQL translator and the many varied systems that support the standard SQL grammar.

25

Agents are programs that respond to search and retrieval operations targeted against a particular data source. The agent loads the appropriate link and translator modules when a request against a specific data source type is first made. The agent then calls upon these module libraries for translation (XLAT), connect, disconnect, search, and retrieval operations.

30

For search operations, an agent will sort the results set by relevance rank and then truncate the set to the maximum number of hits, as specified in the original search

request. This sorting and subsequent pruning of the list of hits is an important precursor to aggregation, which will be discussed shortly.

Agents can reside on the same machine as the data source (recommended) or use a data source's remote APIs for access. More than one copy of an agent can run on a single computer to handle concurrent search and retrieval requests. An agent can be dedicated to a single data source, a group of sources of a particular type, or a range of sources that have a mixture of link requirements.

5

10 Brokers are intermediary components that exist between the requestors of service and the agents that actually perform the service through the backend. They function as special purpose resource coordinators designed to manage the multitude of searches generated from a single request - as caused by a category search for example.

15 A broker typically performs the following tasks:

Validates the request.

20 Expands categories to obtain a list of the data sources available to the application and resolves the source addresses. (Label 1)

Distributes queries to agents for efficient, parallel searching. (Label 2)

25 Aggregates and optionally sorts search results that are returned by the various agents into a single search result set. (Label 3)

Caches search results for subsequent paging operations. (Label 4)

30 Issues requests to agents to retrieve source documents for the user (note that in most cases, the web browser uses the URL returned in the results list to retrieve the document).

Honors timeouts and response options.

5 The degree of responsiveness can vary dramatically from a large set of backend systems contributing to a single request. Some data management systems respond faster than others, and some not at all - possibly due to out of service conditions. To account for this situation, brokers were designed to communicate asynchronously with their agents.

10 To support performance and scalability, an extended search system may contain multiple brokers. This ability to establish a hierarchy of brokers, along with the ability to set up agents co-resident with the sources they support or to dedicate agents to particular sources or types of sources, provides an extended search system, such as IBM Lotus Extended Search, with endless flexibility with regard to changing and expanding environments. Under a multiple broker schema, sources get partitioned across all of the 15 brokers, a design that prevents any one broker from being overwhelmed.

20 In a single broker environment, a search that targets six dozen sources would result in 72 queries being sent to the remote machines and 72 sets of search results being returned to the broker. If each result set contains the maximum number of results, most of the data will be discarded when the broker consolidates and aggregates the data for the list being returned to the requestor (the broker prunes the results and keeps only the top items, up to the maximum number allowed by the search application).

25 With multiple brokers, an entry broker sends a single message to brokers on remote machines. The remote brokers then split the message into multiple requests for the sources (fronted by agents) on their respective machines. Instead of all result sets being returned to one broker, each broker consolidates, aggregates, and prunes the results returned by its agents, and then returns just a single list - containing the top hits - to the entry broker. The entry broker only needs to create a final results set from its own local 30 sources and the consolidated lists returned by the remote brokers.

The broker obtains information about the resources it is to manage from the extended search application's configuration database. This database contains information about data sources and how they should be searched. It also stores network addresses, saved queries, saved search results, and data that was downloaded by a web crawler.

5

An end user can easily update information about a network topology, data sources, and search applications by using an intuitive Administration interface. This interface also provides the gateway through which an end user can run discovery (discussed below), view error message and event data, schedule queries, and work with saved queries and

10 search results.

While the invention has been described with respect to certain preferred embodiments and exemplifications, it is not intended to limit the scope of the invention thereby, but solely by the claims appended hereto.

CLAIMS

1. A method of searching unstructured data comprising:

- 5 a. searching the unstructured data;
- b. returning search results having attributes for aggregation;
- c. aggregating the search results;
- 10 d. returning the aggregated search results to a wrapper; and
- e. entering search result attributes in a nickname table, said nickname table being searchable by a structured data search engine.

15

2. The method of claim 1 comprising searching the unstructured data with an unstructured data search system comprising an extended search broker and a search agent, said method comprising searching the unstructured data with the search agent, 20 aggregating search results obtained thereby with the extended search broker.

3. The method of claim 2 comprising returning search results from the search agent to the extended search broker for aggregation and transferring the aggregated search results and search result attributes to a wrapper.

25

4. The method of claim 3 comprising making the search result attributes available as at least one column in the nickname table.

30

5. The method of claim 1 wherein the nickname table is a relational database nickname table.

6. The method of claim 5 wherein the nickname table is a relational database management system table, the structured data search engine is a relational database management system, and the nickname table is searchable by the relational database management system using a structured query language.

5

7. The method of claim 6 wherein the structured query language is SQL.

8. A method of searching unstructured data comprising:

10 a. searching the unstructured data with a structured data search engine, an unstructured data search agent, and an extended search broker there between;

15 b. returning search results having search result attributes from the search agent to the extended search broker for aggregation; and

c. entering the search result attributes in a nickname table, said nickname table being searchable by a relational database management system using a structured query language.

20

9. The method of claim 8 wherein said extended search broker enters the search result attributes into the nickname table.

25 10. The method of claim 9 comprising entering the search result attributes as a column in the nickname table.

11. The method of claim 9 wherein the structured data search engine queries the search result attributes in the nickname table.

30 12. The method of claim 9 wherein the structured query language is SQL.

13. A computer system comprising an unstructured data search agent, a structured data search engine, and a nickname table, wherein

- a. said computer system is adapted to initiate a search of unstructured data from the structured data search engine through the unstructured data search agent;
- b. said unstructured data search agent is adapted to receive search results having search result attributes from the unstructured data and aggregate the search results;
- c. said computer system is adapted to enter the search result attributes in a nickname table; and
- 15 d. said structured data search engine is adapted to search attributes in the nickname table.

14. The computer system of claim 13 wherein the structured data search engine is a relational data base management system search engine.

20 15. The computer system of claim 14 wherein the relational data base management system search engine is a federated search engine.

25 16. The computer system of claim 15 wherein the unstructured data search agent includes an extended search broker.

17. The computer system of claim 13 wherein the search result attributes are entered into the nickname table as columns in the nickname table.

18. A computer system comprising an unstructured data search system, a structured data search engine, and a nickname table, wherein

- a. said computer system is adapted to initiate a search of unstructured data from the structured data search engine through the unstructured data search system;
- b. said unstructured data search system comprises an extended search broker and an unstructured data search agent; said unstructured data search system is adapted to receive search results with search result attributes from the search agent and return the search results with search result attributes to the extended search broker for aggregation and to aggregate the search results and search result attributes;
- c. said computer system is adapted to enter the search result attributes in the nickname table; and
- d. said structured data search engine is adapted to search the search result attributes in the nickname table.

19. The computer system of claim 18 wherein the structured data search engine is a relational data base management system search engine.

20. The computer system of claim 19 wherein the relational data base management system search engine is a federated search engine.

21. The computer system of claim 18 wherein the unstructured data search system includes an extended search broker.

22. The computer system of claim 18 wherein the search result attributes are entered into the nickname table as a column therein.

23. A computer program product comprising computer readable code to program and 5 configure a computer system to:

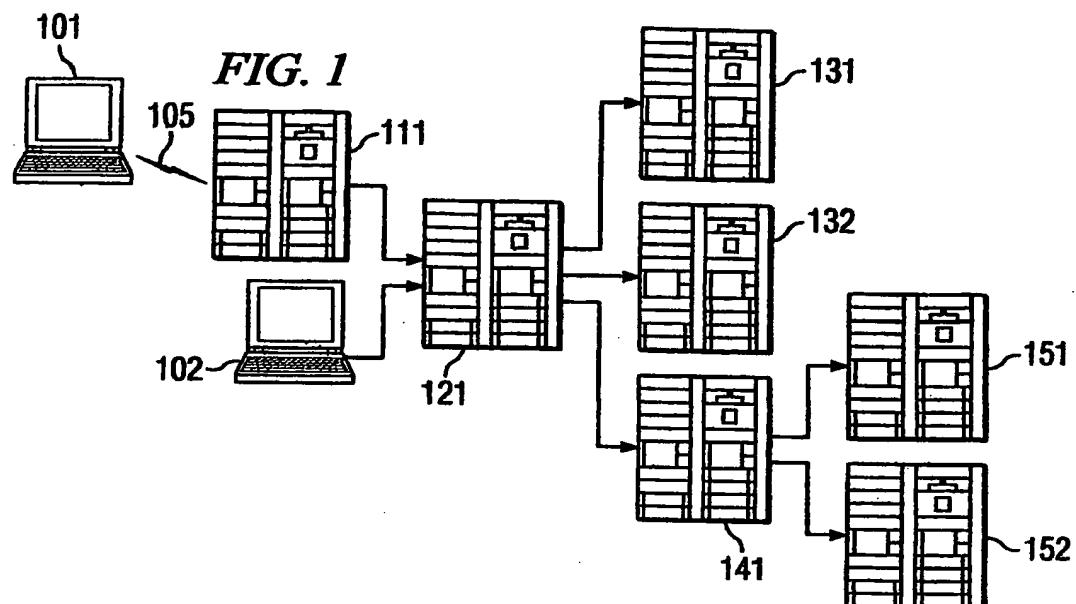
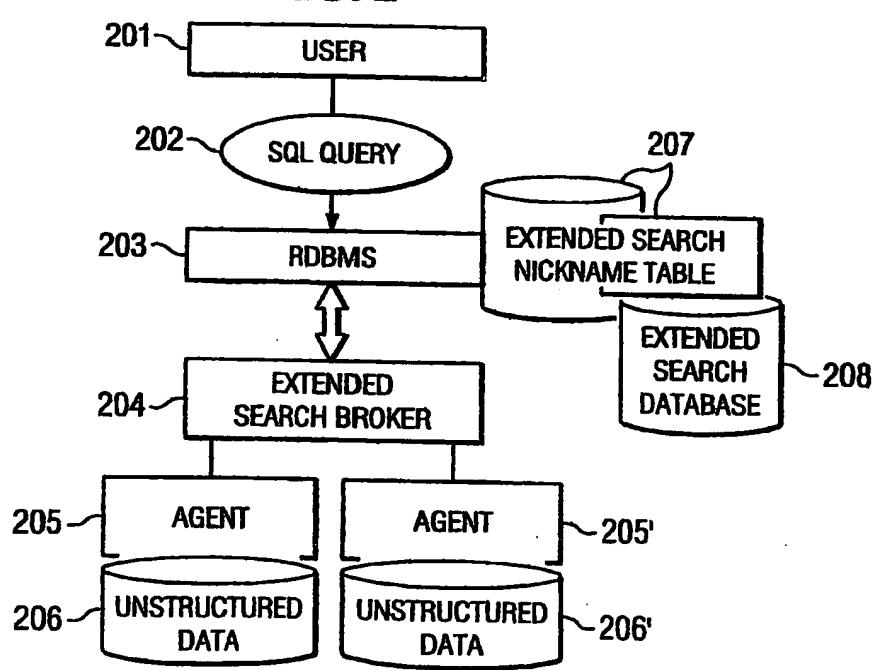
- a. search unstructured data through an extended search broker and a search agent for unstructured data searching;
- 10 b. return search results and search result attributes from the search agent to the extended search broker for aggregation;
- c. aggregate the search results and search result attributes; and
- 15 d. enter the search result attributes in a nickname table, said nickname table being searchable by a relational database management system using a structured query language, whereby to search unstructured data.

24. The computer program product of claim 23 wherein the search results in the 20 nickname table are structured data.

25. The program product of claim 24 wherein the search result attributes are columns in the nickname table.

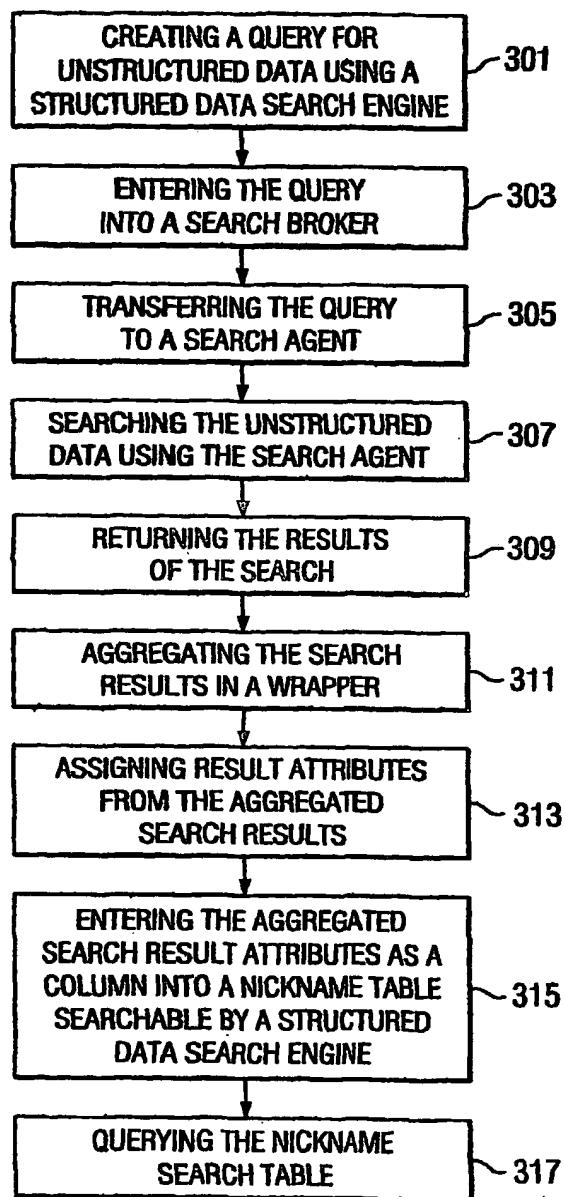
25 26. The computer program product of claim 23 wherein the structured query language is SQL.

1/2

**FIG. 2**

2/2

FIG. 3



INTERNATIONAL SEARCH REPORT

International Application No
PCT/EP2004/002380

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 826 258 A (HARINARIYAN VENKY ET AL) 20 October 1998 (1998-10-20) column 2, line 16 - column 2, line 47 column 4, line 22 - column 6, line 17 column 8, line 31 - column 9, line 23	1-26
X	US 6 282 537 B1 (MADNICK STUART E ET AL) 28 August 2001 (2001-08-28) column 4, line 20 - line 23 column 5, line 42 - line 63 column 9, line 19 - column 10, line 65 column 11, line 62 - column 12, line 62 column 14, line 49 - line 59	1-26 -/-

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the International filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *G* document member of the same patent family

Date of the actual completion of the international search

11 August 2004

Date of mailing of the international search report

06/09/2004

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel: (+31-70) 340-2040, Tx. 31 651 epo nl.
Fax: (+31-70) 340-3016

Authorized officer

DE CASTRO PALOMARES

INTERNATIONAL SEARCH REPORT

International Application No
PCT/EP2004/002380

C(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>ELLEN SPERTUS, LYNN ANDREA STEIN: "Just-in-time databases and the World-Wide Web" CONFERENCE ON INFORMATION AND KNOWLEDGE MANAGEMENT, 'Online! 1998, pages 30-37, XP002292006 PROCEEDINGS OF THE SEVENTH INTERNATIONAL CONFERENCE ON INFORMATION AND KNOWLEDGE MANAGEMENT Retrieved from the Internet: URL:http://delivery.acm.org/10.1145/290000/288638/p30-spertus.pdf?key1=288638&key2=4255122901&coll=Portal&dl=GUIDE&CFID=25625200&CFTOKEN=21019581 > 'retrieved on 2004-08-09! the whole document</p>	1-26
A	<p>GUAN T ET AL: "STRUCTURE-BASED QUERIES OVER THE WORLD WIDE WEB" CONCEPTUAL MODELING-ER. INTERNATIONAL CONFERENCE ON CONCEPTUAL MODELING, XX, XX, 16 November 1998 (1998-11-16), pages 107-120, XP001130882 the whole document</p>	1-26
A	<p>DANIELA FLORESCU, ALON LEVY, ALBERTO MENDELZON: "Database techniques for the World Wide Web: A Survey" ACM SIGMOD RECORD, 'Online! vol. 23, September 1998 (1998-09), pages 59-74, XP002292007 ACM PRESS NEW YORK, NY, USA Retrieved from the Internet: URL:http://delivery.acm.org/10.1145/300000/290605/P059.pdf?key1=290605&key2=7885122901&coll=Portal&dl=GUIDE&CFID=25625200&CFTOKEN=21019581 > 'retrieved on 2004-08-09! the whole document</p>	1-26

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No
PCT/EP2004/002380

Patent document cited in search report		Publication date		Patent family member(s)		Publication date
US 5826258	A	20-10-1998	JP	10222539 A		21-08-1998
US 6282537	B1	28-08-2001	US	5913214 A		15-06-1999
			US	5953716 A		14-09-1999
			AU	3218497 A		05-01-1998
			WO	9745800 A1		04-12-1997